The AI Readiness Gap

What the evidence reveals about schools' readiness for accelerated AI adoption

Introduction

Artificial intelligence has moved into education faster than most schools have been able to respond to it. Today students are using generative AI for homework, for research, and increasingly as something closer to a companion. Teachers are using it to plan lessons and mark work. Safeguarding and IT teams are left retrofitting policies and controls around systems that were designed before generative AI arrived, while trying to plan for what comes next.

The United Kingdom (UK) and the United States (US) educate children in different systems, under different laws, different curricula with different cultural expectations. Despite those differences, the picture in both countries is strikingly similar. AI adoption is high and governance has not kept pace.

This is not a story about AI's role in education. That question has already been settled by the scale of everyday use. The more urgent question is whether schools are equipped to manage it responsibly, and the evidence suggests that for many, the answer is not yet.

1. The Gap, Stated Plainly

In schools across both the UK and the US, the adoption of artificial intelligence has outpaced the ability to govern it.

The scale of the disconnect shows up clearly in UK survey data. A survey of more than 10,000 teachers published in February 2026 found:

- 76% are already using AI in their day-to-day work.[1]
- 49% of schools have no AI policy in place at all.[1]
- 66% have no policy specific to how students use it.[1]
- 76% of teachers do not feel confident advising pupils on appropriate AI use.[2]
- Nearly one in five do not know whether generative AI is even permitted at their school.[2]

The same disconnect appears in US data from 2026. A survey of more than 600 district technology leaders found:

- AI governance and cybersecurity now rank among the top priorities for school IT departments.[3]
- 21% of districts still have no AI guidelines in place.[3]
- Four in five students already use AI for schoolwork.[4]
- Only 6% of teachers say their school's AI policy is clear.[4]

The UK and the US differ in curriculum, culture, and how each system approaches teaching and learning. The pattern in the data is the same in both. The pace of AI adoption is outstripping the training, staffing and infrastructure schools have available to respond to it.

2. The Safeguarding Dimension of the AI Readiness Gap

A missing AI policy is not only a gap in good teaching practice. It is a gap in a school's ability to safeguard its students.

Generative AI tools have already been used to bully, harass, groom, extort and mislead children, and have been linked to the production of AI-generated child sexual abuse material.[5] The Keeping Children Safe in Education 2026 draft guidance explicitly names AI-generated content, including deepfakes, as a safeguarding risk requiring active monitoring.[6] In the US, the Department of Education's national AI guidance identifies human oversight of AI-related decisions affecting students as a top policy priority.[7]

Both frameworks share the same underlying requirement: someone accountable, trained, and positioned to review AI activity and act on it. Without an AI policy, that accountability has nowhere to sit. Without training, staff cannot reliably recognise what they are looking at. Without monitoring visibility, a safeguarding concern involving AI can go unnoticed until it has already escalated.

This is what separates the AI readiness gap from an ordinary technology adoption lag. In most areas of school life, slow adoption of a new tool is an efficiency problem. In safeguarding, a gap between what is happening and what a school can see or respond to is a risk to a child.

3. Where the Gap Shows Up

The readiness gap is not one problem. It shows up differently depending on where in a school's operations you look. Four key areas stand out.

3.1 Policy

The most fundamental gap is the absence of a written policy at all. In the UK, close to half of schools have no AI policy in place, and two-thirds have nothing specific to how students use AI.[1] In the US, one in five districts report having no AI guidelines whatsoever, even as AI governance has become a stated priority for school technology leaders.[3]

A missing policy is not simply an absence of guidance. It means no defined threshold for what counts as a concern, no clarity on who is responsible for spotting one, and no consistent basis for staff to act when something goes wrong. When an AI-related safeguarding issue does arise, a school without a policy is deciding how to respond in the moment, under pressure, rather than following a process agreed in advance.

3.2 Training and Confidence

Even where a policy exists, staff confidence in applying it is low. Three in four UK teachers report not feeling confident advising pupils on appropriate AI use.[2] In the US, only 6% of teachers describe their school's AI policy as clear. [4]

In practice, a policy that staff cannot confidently interpret or apply in the classroom is no policy at all.

3.3 Monitoring and Visibility

Knowing that a student has accessed an AI tool is different from understanding how it was used. Research into student AI use found that a significant proportion of young people are using generative AI weekly, with many copying AI-generated content directly into schoolwork without review.[8] Schools without visibility into the substance of AI interactions, not just whether a site was visited, are working from an incomplete picture of what is actually happening.

In the UK, KCSIE 2026 reflects this shift by introducing an explicit requirement for schools to review the effectiveness of their filtering and monitoring systems annually and to keep a record of those checks.[6]

3.4 Accountability

Every framework examined points to the same underlying principle: a named, trained person must be responsible for reviewing AI-related activity and deciding what action, if any, is required. The US Department of Education identifies human oversight as a top policy priority for AI in education.[7] UK guidance requires designated safeguarding leads to review reports and maintain documented processes for handling them.[6]

Where responsibility is not clearly assigned, an alert or concern can go unreviewed, not because no one cared, but because no one was formally accountable for catching it.

4. What Closing the Gap Requires

Closing the readiness gap does not require new technology as a starting point. It requires four decisions that any school can make regardless of what tools are already in place.

1. **Adopt a written AI policy that names what is and is not permitted.** The policy should cover both staff and student use, distinguish approved from unapproved tools, and be reviewed on a set schedule rather than left static.
2. **Train staff to apply the policy with confidence, not just awareness of it.** Awareness that a policy exists is not the same as knowing how to act on it in a specific situation. Training should cover recognising AI-related risk, not only the rules themselves.
3. **Monitor the substance of AI activity, not only access to it.** Knowing a student visited an AI tool is not sufficient. Effective monitoring captures what was asked and what was returned, giving reviewers the context needed to assess whether a concern is genuine. KCSIE 2026 formalises this expectation by requiring an annual review of filtering and monitoring effectiveness, with a record kept of the outcome.[6]
4. **Assign clear, named accountability for reviewing flags and deciding what happens next.** A designated, trained person, with defined backup cover for absence, should be responsible for reviewing AI-related concerns and determining the appropriate response. Both UK and US frameworks treat this as a baseline requirement, not best practice.[6][7]

Schools do not need to solve all four at once, and few will start from the same position. What matters is having a clear view of where the gaps sit today, and a plan for closing them in order of risk. Whether that work happens through internal policy development, staff training, or evaluating external safeguarding tools, these four decisions are the standard against which any approach should be measured.

Source Validation

[1] National Education Union survey, February 2026. A survey of more than 10,000 teachers across England, Wales and Northern Ireland found that 76% are already using AI in their day-to-day work, while 49% of schools have no AI policy in place at all and 66% have no policy specific to how students use AI.

[2] DfE-commissioned research, cited in the Tony Blair Institute's January 2026 paper. The research found that 76% of teachers do not feel confident advising pupils on appropriate AI use, and nearly one in five do not know whether generative AI is permitted at their own school.

[3] CoSN, State of EdTech District Leadership Survey 2026. A survey of more than 600 district technology leaders found that AI governance and cybersecurity now rank among the top priorities for school IT departments, while 21% of districts report having no AI guidelines in place.

[4] Stanford HAI, AI Index Report 2026. The report found that four in five US students already use AI for schoolwork, while only 6% of teachers describe their school's AI policy as clear.

[5] NSPCC. The NSPCC has documented cases of generative AI being used to bully, harass, groom, extort and mislead children, alongside the broader risk of AI-generated child sexual abuse material.

[6] Keeping Children Safe in Education 2026, draft consultation, Department for Education, published 12 February 2026. Paragraph 21 identifies AI-generated content, including deepfakes, as a safeguarding risk. Paragraphs 166 to 169 introduce a requirement for schools to review the effectiveness of filtering and monitoring systems annually and to keep a record of those reviews.

[7] US Department of Education. National guidance on AI in education identifies establishing human oversight, sometimes described as human-in-the-loop, as a top policy priority for AI systems used in educational settings.

[8] National Literacy Trust, 2025. Research based on a survey of more than 60,000 young people aged 13 to 18 found that two in three reported using generative AI, with almost half using it at least weekly, and one in four admitting to copying AI-generated content directly into homework without editing it.